At Digital Equipment Corp., a system called XCON uses knowledge of the company's extensive product line to check—and correct—orders for large computers. At Filene's, a Boston department store, executives obtain information about the company's finances, employees, and merchandise from Intellect—a computer program that understands questions typed in everyday English. At Machine Intelligence Corp. (Sunnyvale, Calif.), a computer-based vision system enables a robot to locate tools scattered on a tabletop.

These systems are harbingers of a new wave of smart computers just beginning to emerge from computer science laboratories. Able to mimic—if not duplicate—human thought processes, such as reasoning, perception, and even learning, these artificial intelligence (AI) systems are expected to open vast new opportunities for automation in the office, factory, and home. In the process, many observers believe, they will profoundly alter the way people work, live, and think about themselves.

Key to the operation of AI systems is a radically new style of computer programming. Conventional programs instruct a computer to solve a problem by following a rigidly predetermined sequence of steps called an algorithm. But many problems are too complex to fit this algorithmic straightjacket; step-by-step solutions are either unknown or too inefficient even for a fantastically fast machine like a modern digital computer.

As a result, AI systems use knowledge about a problem to suggest shortcut solutions—a technique known as heuristic problem-solving. Because knowledge is key to this technique, AI researchers have developed—and continue to develop—methods for efficiently representing and processing facts and ideas on a computer. These techniques vary sharply, but they are all basically formalizations of natural language. Some systems use *if–then* rules, such as "if a cat is large and has stripes, then it's a tiger." Another common technique employs pyramid-shaped networks of related ideas ("semantic networks") to deduce facts, such as "all cats eat meat, but only tigers have stripes." Some systems apply formal logic to deduce facts from a set of elementary propositions called axioms.

When solving a problem heuristically, AI systems do not follow a prescribed sequence of steps. Instead, they are programmed to follow general problem-solving procedures, such as breaking the problem down into easier problems or using information about the problem to suggest and test hypothetical solutions. These broad guidelines are then used to determine a specific sequence of processing steps.

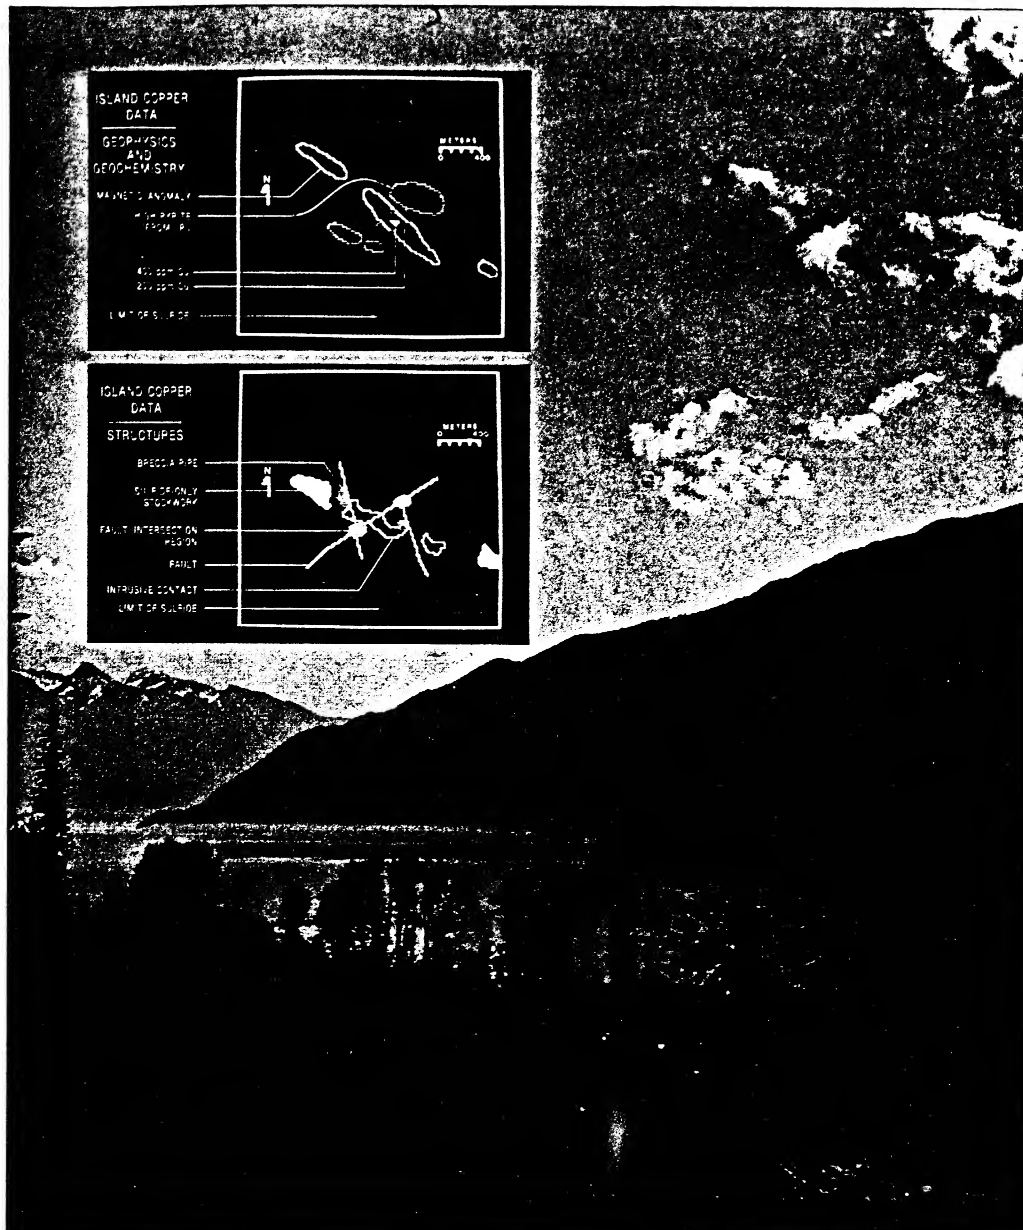

To solve a problem in medical diagnosis, for example, a system may request some initial information about a patient's condition from an attending physician. It then searches its memory for a rule that fits the initial data. This rule may supply an immediate diagnosis, but usually it triggers a request for additional information. This process continues until the system reaches a conclusion. At any time, a user may ask the system to explain a line of questioning or the relevance of a symptom.

An AI system's knowledge base usually comprises many rules and facts. For example, the XCON system developed by DEC and Carnegie-Mellon University uses some 1500 rules and 500 product descriptions to determine permissible combinations ("configurations") of system components. In checking a computer order, XCON typically runs through 1000 rules—a process that takes about two minutes.

AI researchers have dreamed of building thinking machines almost since the invention of the computer itself. But early failures in machine translation and computer chess, compounded by frequently overblown claims, gave AI a bad name—especially in industry. "The history of AI has been fraught with false expectations," says one researcher.

A few universities and private think tanks kept the flame alive—largely with the help of Defense Dept. funding. The leading research centers—Stanford, MIT, and Carnegie-Mellon University—developed different approaches to

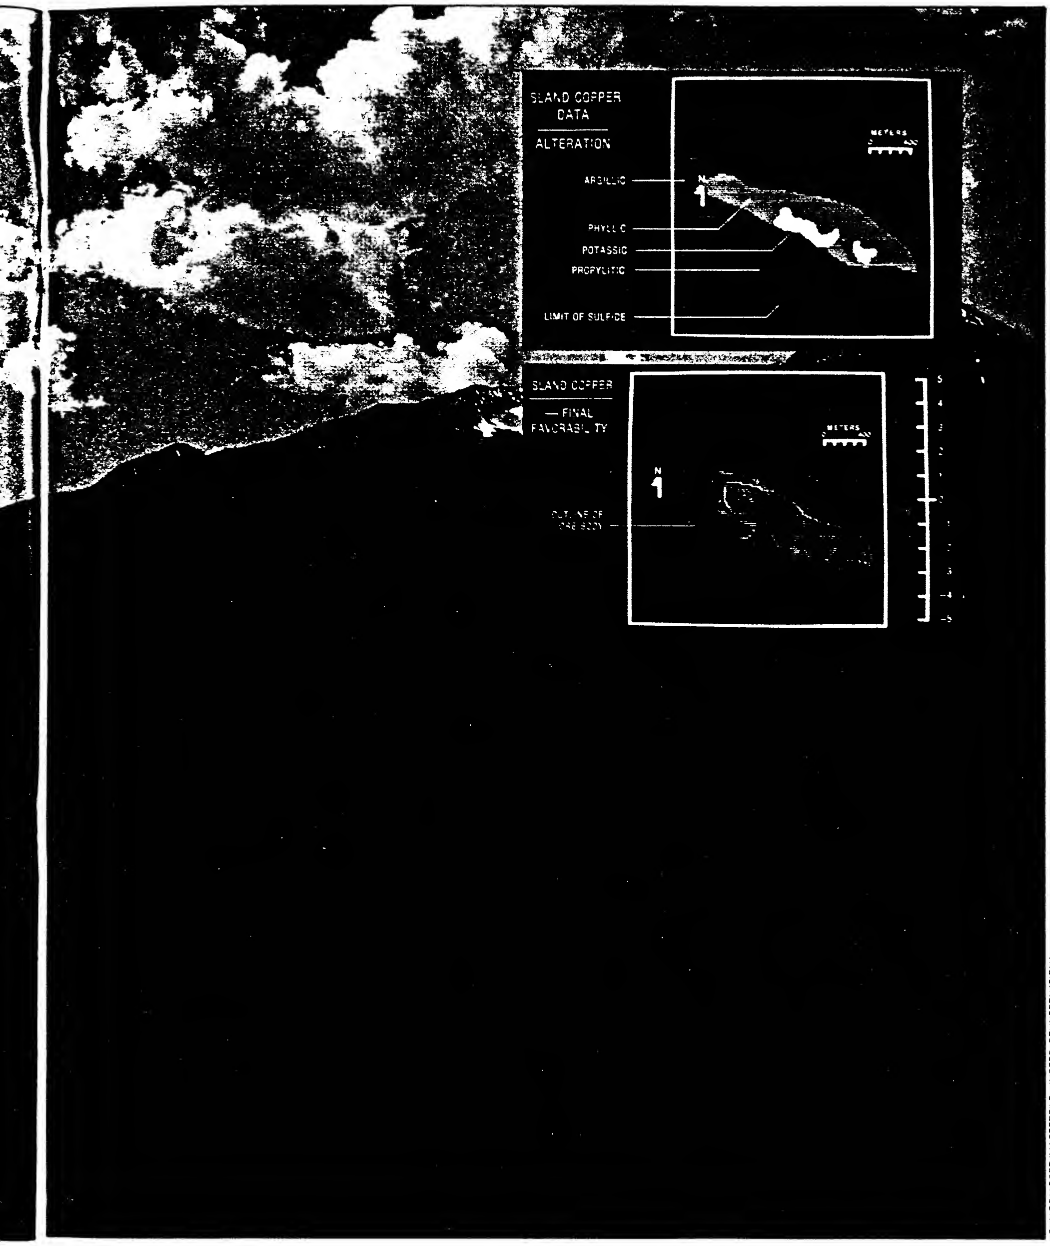

JOHN DE VISSER/MASTERFILE INSETS: SRI INTERNATIONAL

The Prospector expert system predicts mineral ore deposits from geological data entered as maps or through a dialogue with a user. The maps on the left depict the geology of the Island Copper deposit in northwestern Vancouver Island, British Colombia. The map on the right shows predicted deposit location (orange equals high probability of deposit; blue, low probability) versus actual location of ore body (outline).

ARTIFICIAL INTELLIGENCE: MAKING COMPUTERS SMARTER

by Paul Kinnucan

A new wave of supersmart computers is about to invade the office and factory

AI that frequently sparked fierce debates among their adherents. Even today AI researchers are ready to argue the merits of various knowledge representation schemes and whether machine intelligence should be patterned after human cognition. But all agree that computers need access to large pools of knowledge to solve many practical problems.

Now the semiconductor revolution is giving AI researchers a chance to emerge from their ivory tower. AI systems that once required a roomful of equipment now fit in a cabinet the size of an office refrigerator. The reduction in prices for AI systems has been equally dramatic—from several million dollars a few years ago to as little as $50,000 today. As a result, industry is beginning to view AI with renewed interest—and respect.

Signs of a serious corporate interest in AI abound. IBM, Texas Instruments, and Xerox have harbored AI research groups for years. Now other technically savvy companies—Hewlett-Packard, Atari, and Schlumberger, to name a few—are rushing to set up their own centers. Companies, many of them started by prominent academic researchers, are springing up to commercialize AI technology.

There is also intense interest in AI abroad—a development that may signal an end to U.S. domination of the field. Japan, in particular, plans to spend $450 million over the next decade to develop an advanced computer targeted specifically at AI applications. For example, the system's design calls for development of knowledge processors capable of storing and retrieving as many as 20,000 rules and 100 million data items. This "fifth generation" computer is the centerpiece of Japan's plan to attain world leadership in AI technology by the end of the century.

Most observers believe the initial commercial applications for AI technology will emerge in three areas: computer-based consultants that assist users by offering advice and rendering judgments in fields that require specialized knowledge; electronic information systems that understand ordinary English and hence are accessible to nontechnical users; and artificial vision systems

aimed at industrial applications, such as smart robots and automated inspection, or military applications such as automatic photointerpretation.

Expert systems represent the leading edge of AI technology; more than 50 have been built thus far. AI researchers build expert systems by picking a human specialist's brain and encoding the result as *if–then* rules—a procedure known as knowledge engineering.

Debriefing an expert may take months or years, but the payoff can be significant. An expert system called ONCASYN, for example, helps doctors keep track of drug therapy for victims of Hodgkin's disease at a Stanford University cancer clinic. Drug protocols for cancer patients are so complex that it's easy for doctors to make mistakes. Other medical systems developed at Stanford include MYCIN, a 500-rule system that diagnoses blood diseases, and PUFF, a system that diagnoses lung ailments.

Expert systems have also demonstrated their worth in industrial applications. For example, researchers at most of the major pharmaceutical firms have been using DENDRAL—another Stanford-developed expert system—for more than a decade to identify organic molecules from mass spectrograms and other lab data. The system excels chemistry PhDs in this task because it can consider more possible identities for a mystery molecule, according to Edward Feigenbaum, who developed DENDRAL with Nobel laureate Joshua Lederberg in the mid-sixties. Feigenbaum, now chairman of Stanford's computer science department, has started two companies—Teknowledge Inc. and Intelligenetics Inc.—to market expert systems.

To shorten the development cycle for expert systems, Stanford has created a skeleton system called EMYCIN—basically MYCIN without its medical knowledge. IBM has used EMYCIN to build a prototype expert system for diagnosing malfunctions in computer disk drives. UNITS, another system-building tool developed at Stanford, enables knowledge engineers to build semantic networks consisting of frames—knowledge packages that describe both the fixed and variable attributes of objects. Stanford, Carnegie-Mellon, and other research centers are also developing knowledge representation languages (KRLs) aimed at reducing the cost of building expert systems.

To solve many problems, AI researchers believe, a computer must be capable of fuzzy thinking—developing a line of reasoning based on uncertain or partial evidence. Otherwise computers will be restricted to problems that admit only a true or false answer rather than a spectrum of answers. Many expert systems already demonstrate a fuzzy thinking capability. For example, Prospector—developed by SRI International (Menlo Park, Calif.) for identifying mineral deposits—allows users to specify their "confidence level" when answering the system's requests for information about a site. Prospector's rules also rank evidence for a deposit according to its sufficiency and the necessity for establishing the existence of the deposit. At the end of a consultation, Prospector gives the probability of ore being located at a site at a given concentration—a more meaningful reply than a simple yes/no answer. The system reportedly identified a molybdenum deposit worth several million dollars last year in Canada.

Many observers believe that expert systems will eventually put expert legal, financial, and medical advice at the fingertips of anyone with access to a personal computer. Cognitive Systems Inc., a firm recently founded by Roger Schank, chairman of Yale University's computer science department, is devel-

How AI systems represent knowledge

AI systems, like people, employ language—strings of symbols—to represent facts and ideas in their memories. Indeed, the knowledge representation languages (KRLs) used by AI systems are essentially highly simplified versions of everyday language. Although KRLs lack natural language's expressive power and flexibility, they allow knowledge to be stored compactly and processed efficiently—important considerations because of the limited memory capacity and processing power of computers.

Many expert systems represent knowledge by formalized *if–then* rules—"if the patient has a runny nose and a fever, then the patient is likely to have a cold." These production rules suit the antecedent–consequent reasoning often employed by human experts. AI systems mimic such reasoning processes by chaining forward or backward through the rules.

In forward chaining, a system seeks to find a rule whose antecedent matches initial data. The consequent of the applicable rule then creates the condition for firing a new rule. This process continues until the system reaches a conclusion. In a medical system, the initial data, "patient has runny nose and a fever" might fire a rule with the "has cold" consequent. This rule would then fire the rule: "if patient has cold, then prescribe two aspirin and put to bed," thus ending the session. In backward chaining, a system selects a rule and tries to match its antecedent against the initial input data. If this fails, it searches for a rule whose consequent would—if applicable—fire the initial rule. The medical system might select the cold rule and ask the user if the patient has a fever. If the user does not know, the system might find a fever rule, such as "if patient's temperature is greater than 98 degrees Fahrenheit, then patient has a fever." The system would then ask the user for the patient's temperature.

Both types of chaining exemplify general reasoning procedures used by AI systems. Forward chaining is a type of bottom-up (data-driven) reasoning, while backward chaining is a form of top-down (goal- or expectation-driven) reasoning. Many systems combine both types of reasoning in a problem-solving procedure known as generate-and-test. Initial data produce hypotheses, which then guide the search for more information.

Semantic networks are another popular knowledge representation scheme used in AI systems. A semantic network may be represented by a graph: Nodes stand for concepts, such as robin, bird, wing; connnecting arcs symbolize relationships, such as is-a or has-a. For example, a semantic network would show that Bill is a robin by using an is-a arc to join a Bill-node to a robin-node. (Actually, graphs are just a convenience for visualizing semantic networks. In computer memory, semantic networks consist of records linked by pointers; the records, which correspond to nodes, contain "slots" for the pointers. For example, an is-a slot in a robin-record would contain a pointer to a bird-record.)

Semantic networks take advantage of the abstract to concrete relationship among concepts. A semantic network dealing with birds might connect the abstract concepts "physical object" to "animate being" to "bird," which would then be linked to the concepts "robin," "eagle," and so on. This pyramid-like structure conserves memory since every lower object in the hierarchy inherits the properties of higher objects. For example, a robin inherits all the properties of a bird, animate being, and physical object; hence the robin-node does not have to repeat those properties.

oping forerunners of such home systems. For example, the TAD (Tax Advisor) system now being developed by the firm will not only fill out IRS tax forms for a user but will also advise the user on what income must be reported, what expenses are deductible, and other regulations governing income tax reports. The hardware required to run TAD, which interacts with the user in natural English, sells for about $60,000, a price Schank expects to drop to as low as $10,000 within a few years. Schank believes that expert systems for the home will have to await the arrival of a new generation of 32-bit microcomputers. The current 16-bit generation does not have enough speed and memory capacity to power AI software, he says. For this reason, his company is aiming its current systems at brokerage houses, banks, and other large financial institutions. For example, Cognitive Systems is negotiating with H&R Block to install TAD systems in its local offices.

Even as the current generation of rule-based systems is beginning to move out of the laboratory, AI researchers are launching development of a new wave of laboratory systems based on "deep knowledge"—the memories, scientific theories, mental models, and other forms of knowledge out of which experts compile the rules used to drive current systems. AI researchers believe such systems will be more robust than the current rule-based systems. "What scares us about all the hype surrounding expert systems is that they are very brittle," says John Seeley Brown, director of the Cognitive and Instructional Sciences Group at Xerox Corp.'s Palo Alto (Calif.) Research Center (PARC). "As soon as you go over the boundary of what they were designed to do, they collapse into a quivering heap." Human experts, on the other hand, can fall back on their deep knowledge to deal with a new problem, Brown points out.

Deep knowledge will be the key to the performance of an advanced system for diagnosing computer malfunction being developed at Stanford under IBM sponsorship. The new system will incorporate functional models of computer system components that will enable it to explain system malfunctions by proving how they might occur, such as through a component failure. The system will not be limited to the knowledge of the pathology of computer behavior as is the current IBM system, says Michael Genesereth, an assistant professor of computer science at Stanford.

It is anybody's guess when such systems will reach the commercial market. Many research issues need to be resolved, such as what deep knowledge is relevant and whether current knowledge representation schemes are adequate.

Besides serving as computerized consultants, expert systems may fulfill other roles in the coming decades. For example, Bolt Beranek and Newman Inc. in Cambridge, Mass., is developing an expert system intended to teach Navy personnel how to operate shipboard steamplants. Called Steamer, the system simulates a frigate steamplant on its color display, enabling a user to experiment with various operating procedures. In addition, the system incorporates a knowledge base about the frigate steamplant that enables it to explain its operation and terminology.

AI systems are just beginning to demonstrate the pedagogical skills needed to convey skills and knowledge effectively. At Xerox PARC, for example, a computer-aided instruction (CAI) system called WEST engages children in a video game that requires the use of basic arithmetic skills, thereby sweetening the pill of learning. Another program called DEBUGGY diagnoses arith-

Besides conserving storage space, semantic networks also facilitate deductive reasoning. For instance, the property inheritance feature allows a system to deduce that "robins have wings" from the relationships "all birds have wings" and "a robin is a bird."

A recent innovation called frames greatly increases the power of semantic networks. Frames are two-part nodes: A fixed part specifies the permanent features of objects, while a variable part identifies changeable attributes plus the values they may assume. A frame may even specify a default value for changeable attributes. Among the permanent attributes of seagulls, a frame might list the facts that they are birds and live by the water. The frame might include plumage-color as a variable feature having the values brown or gray-and-white, with the latter the default value.

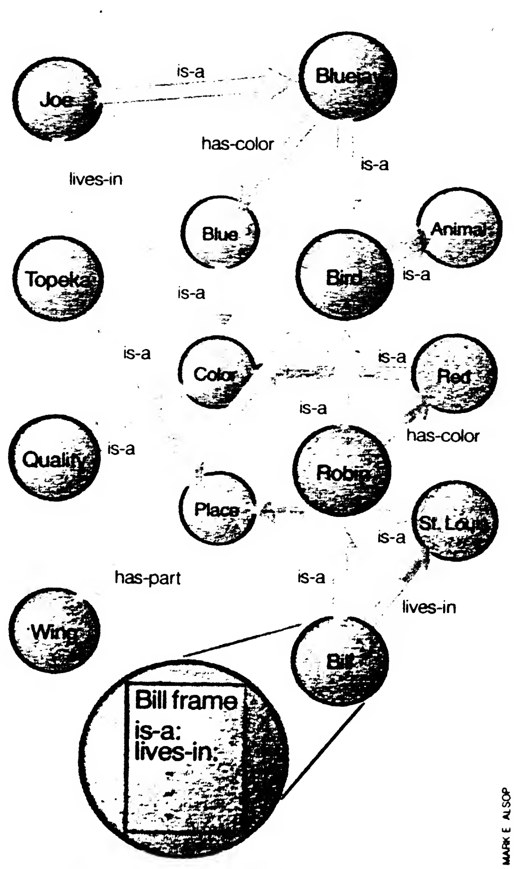

Fragment of a hypothetical semantic network representing knowledge about birds.

Frames allow AI systems to expand their knowledge by fitting random information into a pre-existing conceptual framework. For example, in reading a story about a seagull named Jonathan, the system could add Jonathan to its knowledge base by copying a seagull frame and filling in its variable slots with information garnered about Jonathan—a procedure known as "instantiating" a frame. The plumage-color default value would allow the system to know that the phrase the "gray-and-white bird" refers to Jonathan without being told explicitly that Jonathan is gray and white. Frames also allow a system to form expectations. From the fact that Jonathan is a seagull, the system would expect to find descriptions of the sea in the story.

Some AI systems use the symbolism of formal logic to represent knowledge. This symbolism expresses assertions as truth functions of their component parts. The fact that all birds have wings would be expressed as FOR ALL X, BIRD(X) IMPLIES HAS-WINGS(X). The truth functions, BIRD(X) and HAS-WINGS(X), have the value "true" when X is a bird. Expressing a body of facts in this manner allows new facts to be deduced with mathematical certainty by a set of rules known as the predicate calculus. However, technical problems have prevented formal logic from attaining widespread use.

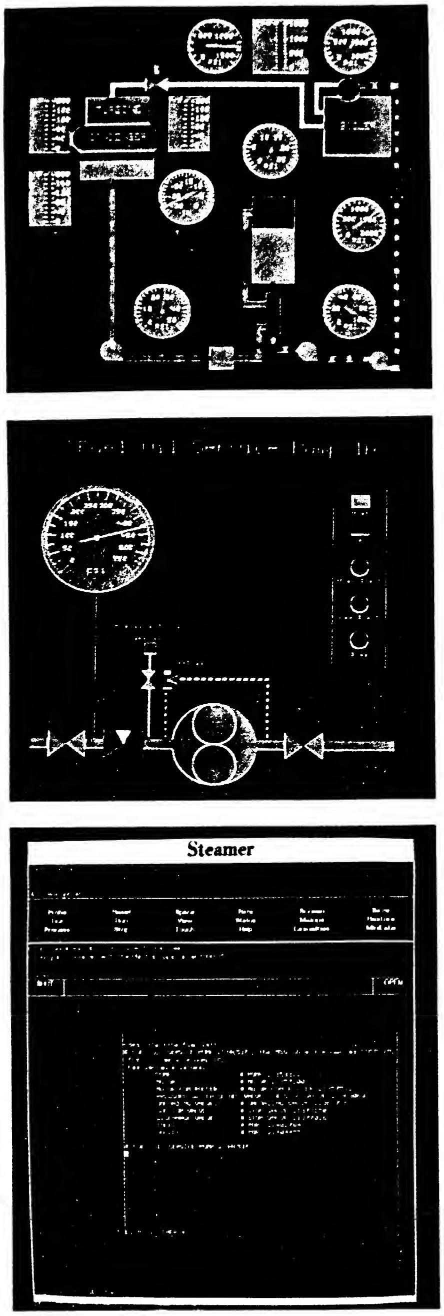

WALTER BIBIKOW

The Steamer computer-aided instruction system may someday help teach U.S. Navy recruits how to operate ship propulsion systems. A mathematical simulator allows users to practice operating procedures and see the results on an animated color display (top). The system's multi-level display capability (middle) shows the steamplant's structure at increasing levels of detail. A knowledge base consisting of frames enables the system to describe the steamplant's components and operating procedures (bottom).

metic mistakes by having a child do a series of subtraction problems designed to reveal faulty ("buggy") calculation procedures. The program excels most teachers in diagnosing subtraction errors, says Xerox's Brown.

In addition to teaching ideas, AI systems have also demonstrated an ability to discover them. A system called Eurisko, for example, last year discovered a strategy for winning the war game, Traveler, that enabled Eurisko's developer, Doug Lenat of Stanford, to win the game's national championship. Eurisko works by a kind of brainstorming process—it combines concepts in its knowledge base to form new ideas, keeping those that look interesting and discarding the rest. Lenat is currently using Eurisko to generate and evaluate logic designs for a new type of multilayer, integrated-circuit chip developed at Stanford.

AI researchers have had a difficult time trying to endow computers with the seemingly simple functions of sight and language. This is because both functions assume a great deal of knowledge and reasoning—much of it unconscious in people. "It is amazing how much inferencing takes place in understanding even the simplest story," observes Yale's Schank.

But commercial systems able to understand ordinary English are beginning to appear. Artificial Intelligence Corp. in Waltham, Mass., for example, has been marketing an English language information system for two years. Called Intellect, the system enables users to store information on a specific subject in a computer database and later retrieve the information via questions typed in English on a keyboard. For example, to retrieve information about sales personnel from an employee file, a user might type "Show me all the salesmen in San Francisco who earn less than $50,000." The system uses linguistic knowledge in the form of a grammar and dictionary to translate such questions into formal queries, such as "DISPLAY FILE=employee FOR JOB-DESCRIPTION=sales .AND. SEX=male .AND. CITY=San Francisco .AND. SALARY<$50,000." The system's formal query language processor then retrieves and displays the requested information on a CRT screen. (The translation software is called a "natural language front-end" because it stands between the user and the database's formal query language processor.) To date, AIC has installed more than 80 Intellec software packages, which sell for about $70,000 and run on large IBM computers. Semantek, a Sunnyvale, Calif., firm started by former SRI International researcher, Gary Hendrix, expects to introduce a database system for personal computers that will have a natural language capability. Although it will not cope with the full range of English syntax because of memory restrictions imposed by a personal computer, the system will sell for less than $2000, says Hendrix.

To allow easy access to the numerous computer databases already in existence, some AI researchers are developing "transportable" natural language front-ends that can adapt to an existing database. SRI International, for example, has built a front-end called TEAM (Transportable English Access Mechanism) that elicits information about the structure and contents of an existing database from its manager during an initial interview. TEAM then uses this information to modify its translation mechanism so as to produce queries in the host system's native query language.

Systems like Intellec and TEAM, which are built by combining a natural language front-end with a database management system (DBMS), perform remarkably well because they take advantage of a DBMS' ability to answer a wide range of questions about a subject without knowing anything about the subject itself. The only thing a DBMS knows about are information retrieval concepts such as files, records, fields, and search, and linguistic concepts such as words and grammatical rules. All the front-end adds is knowledge about natural language. Such systems require programming to answer questions that depend on knowledge of the subject matter ("world knowledge") as opposed to knowledge of the system and its contents ("ego knowledge"). For example, a system might successfully answer the questions "What was IBM's net income last quarter?" and "What were IBM's net expenses last quarter?" and yet fail the question "What is IBM's net profit?" The reason? The system lacked a "profit" entry in its database. The last question reveals that the knowledge of financial analysis apparent in the first questions is only an illusion.

Though unlikely to find commercial application soon, AI systems are beginning to emerge that apply world knowledge to such language understanding tasks as reading and summarizing news stories. For example, IPP, a system developed by Roger Schank's group at Yale, uses knowledge of terrorism to read and summarize news wire stories. Knowledge of typical terrorist incidents stored as standardized "scripts" in its memory enables IPP to deduce information not explicitly mentioned in a story. For example, from a script for terrorist attacks in Israel, the system

may deduce that a terrorist bombing in Jerusalem was caused by Palestinians even though the story failed to mention the nationality of the terrorists. While reading a story, IPP builds up a language independent representation of the story by instantiating a standard script—filling in "slots" in the script left open for specific information, such as the time of an attack and the names of victims. It then generates a summary of the story from the instantiated script.

Besides being able to read stories, IPP has other nifty features. For example, it can produce summaries in other languages, such as Spanish, Hebrew, and Russian, thus providing a translation capability. In addition, IPP has the ability to learn from its reading and apply this new knowledge to increase its ability to understand new stories. For example, in reading stories about terrorist incidents in Ireland, it formed the generalization that such attacks are caused by the Irish Republican Army. When it subsequently read a story that did not mention the IRA, it simply assumed this information and inserted it in its summary.

The careless application of such a generalization could lead IPP to a false conclusion, of course. Indeed, this is a characteristic feature of AI systems that separates them from conventional systems. Standard systems only make mistakes when there are errors in their algorithms or data; their decision-making processes are strictly logical. AI systems, on the other hand, often depend on educated guesswork to reach a conclusion. This enables them to reach conclusions in situations where a conventionally programmed computer would be at a loss. However, conclusions based on guesswork can always be wrong. Hence, AI systems can make mistakes even when working with correct data.

Just as AI researchers have developed computers able to understand users in their native tongue, they are also developing systems with the ability to see. Commercial systems based on AI research are beginning to appear. Several companies, including Machine Intelligence Corp., Automatix Inc. (Burlington, Mass.), and Octek Inc. (Burlington, Mass.), are marketing versions of the Vision Module developed at SRI International in the seventies. These systems, which consist of a microcomputer and TV camera, see—and know—only a world of two-dimensional silhouettes, and hence have severe limitations. They can only recognize objects that contrast sharply with their background since they only recognize an object by its silhouette. And because they lack any conception of a 3-D world, they cannot recognize objects from different perspectives or objects that overlap.

Although viewing conditions can be controlled to skirt these limitations in many industrial applications, AI researchers, such as Thomas O. Binford of Stanford University, argue that more versatile vision systems would enable computers to operate in the vastly greater number of potential applications where viewing conditions cannot be rigidly controlled. As a result, Binford and others are developing systems that apply extensive knowledge of the world to visual analysis. For example, ACRONYM, a vision system developed by Binford, uses 3-D models of passenger jets to predict their appearance in aerial photographs and thereby recognize them from different perspectives. Berthold Horn of MIT and Christopher Brown of the University of Rochester have developed systems that deduce the shape of objects from their shaded appearance in an image—the same trick that people use to distinguish between a disk and a ball in a picture.

Some vision systems are even emerging that use expert system techniques. For example, an experimental radar target classification system being developed for the Navy by Advanced Information and Decision Systems Inc. (AI&DS) uses rules for visual identification of ships to distinguish between different types of destroyers, cruisers, and aircraft carriers.

The AI&DS system uses a combination of bottom-up and top-down processing to extract high-level features of an unknown ship, such as its length and the location of major superstructures and decks, from low-level features of the radar image, such as edges, blobs, and blips (see photo, page 69). If-then rules are used to match the mystery ship's features with those of known ships stored as three-dimensional models in the system's memory. For example, a typical rule states the likelihood that the unknown ship is identical to a known ship with a similar pattern of superstructure peaks. Because of the poor quality of the radar images, the system usually makes several passes before arriving at a positive identification. The system first guesses at an identification from initial features, then attempts to confirm the guess by finding features corresponding to the hypothetical ship's model. This procedure is repeated until a positive identi-

Knowledge-based translation

Yale's MOP-TRANS knowledge-based translation system tries to understand the subject of a news story in Spanish before attempting an English translation. The system's understanding of the story takes the form of a memory packet (right) created by individualizing stereotyped scenes stored in its knowledge base, using details garnered from the Spanish original (top left). The memory packet is then used to generate the English translation (middle left)—actually a paraphrase as a comparison with a human translation (bottom left) shows.

Spanish original:
La policia realiza intensas dilegencias para capturar a un presunto maniático sexual que dio muerte a golpes y a puñaladas a una mujer de 55 años, informaron fuentes allegadas a la investigacion.

Machine translation:
The police are searching for a sex maniac because he killed a 55-year-old woman.

Human translation:
Police are searching for a presumed sex maniac who beat and stabbed to death a 55-year-old woman.

```
BEGIN-MEMORY-PACKET: story #1
  CONCEPT: police-capture
  ACTOR: police #1
  OBJECT: sex-maniac #1
  ARREST-SCENE: scene #1
    CONCEPT: arrest
    ACTOR: police #1
    OBJECT: sex-maniac #1
    RECIPIENT: police #1
    TIME: future #1
    RESULTS-FROM: scene #2
      INITIATED-BY: scene #3
  SEARCH-SCENE: scene #2
    CONCEPT: police-search
    ACTOR: police #1
    OBJECT: sex-maniac #1
    TIME: now #1
  CRIME-SCENE: scene #3
    CONCEPT: crime
    ACTOR: sex-maniac #1
    OBJECT: woman #1
      CONCEPT: human
      GENDER: female
      AGE: 55-years
    RESULT: death #1
    ACT-OF: homicide
END-MEMORY-PACKET: story #1
```

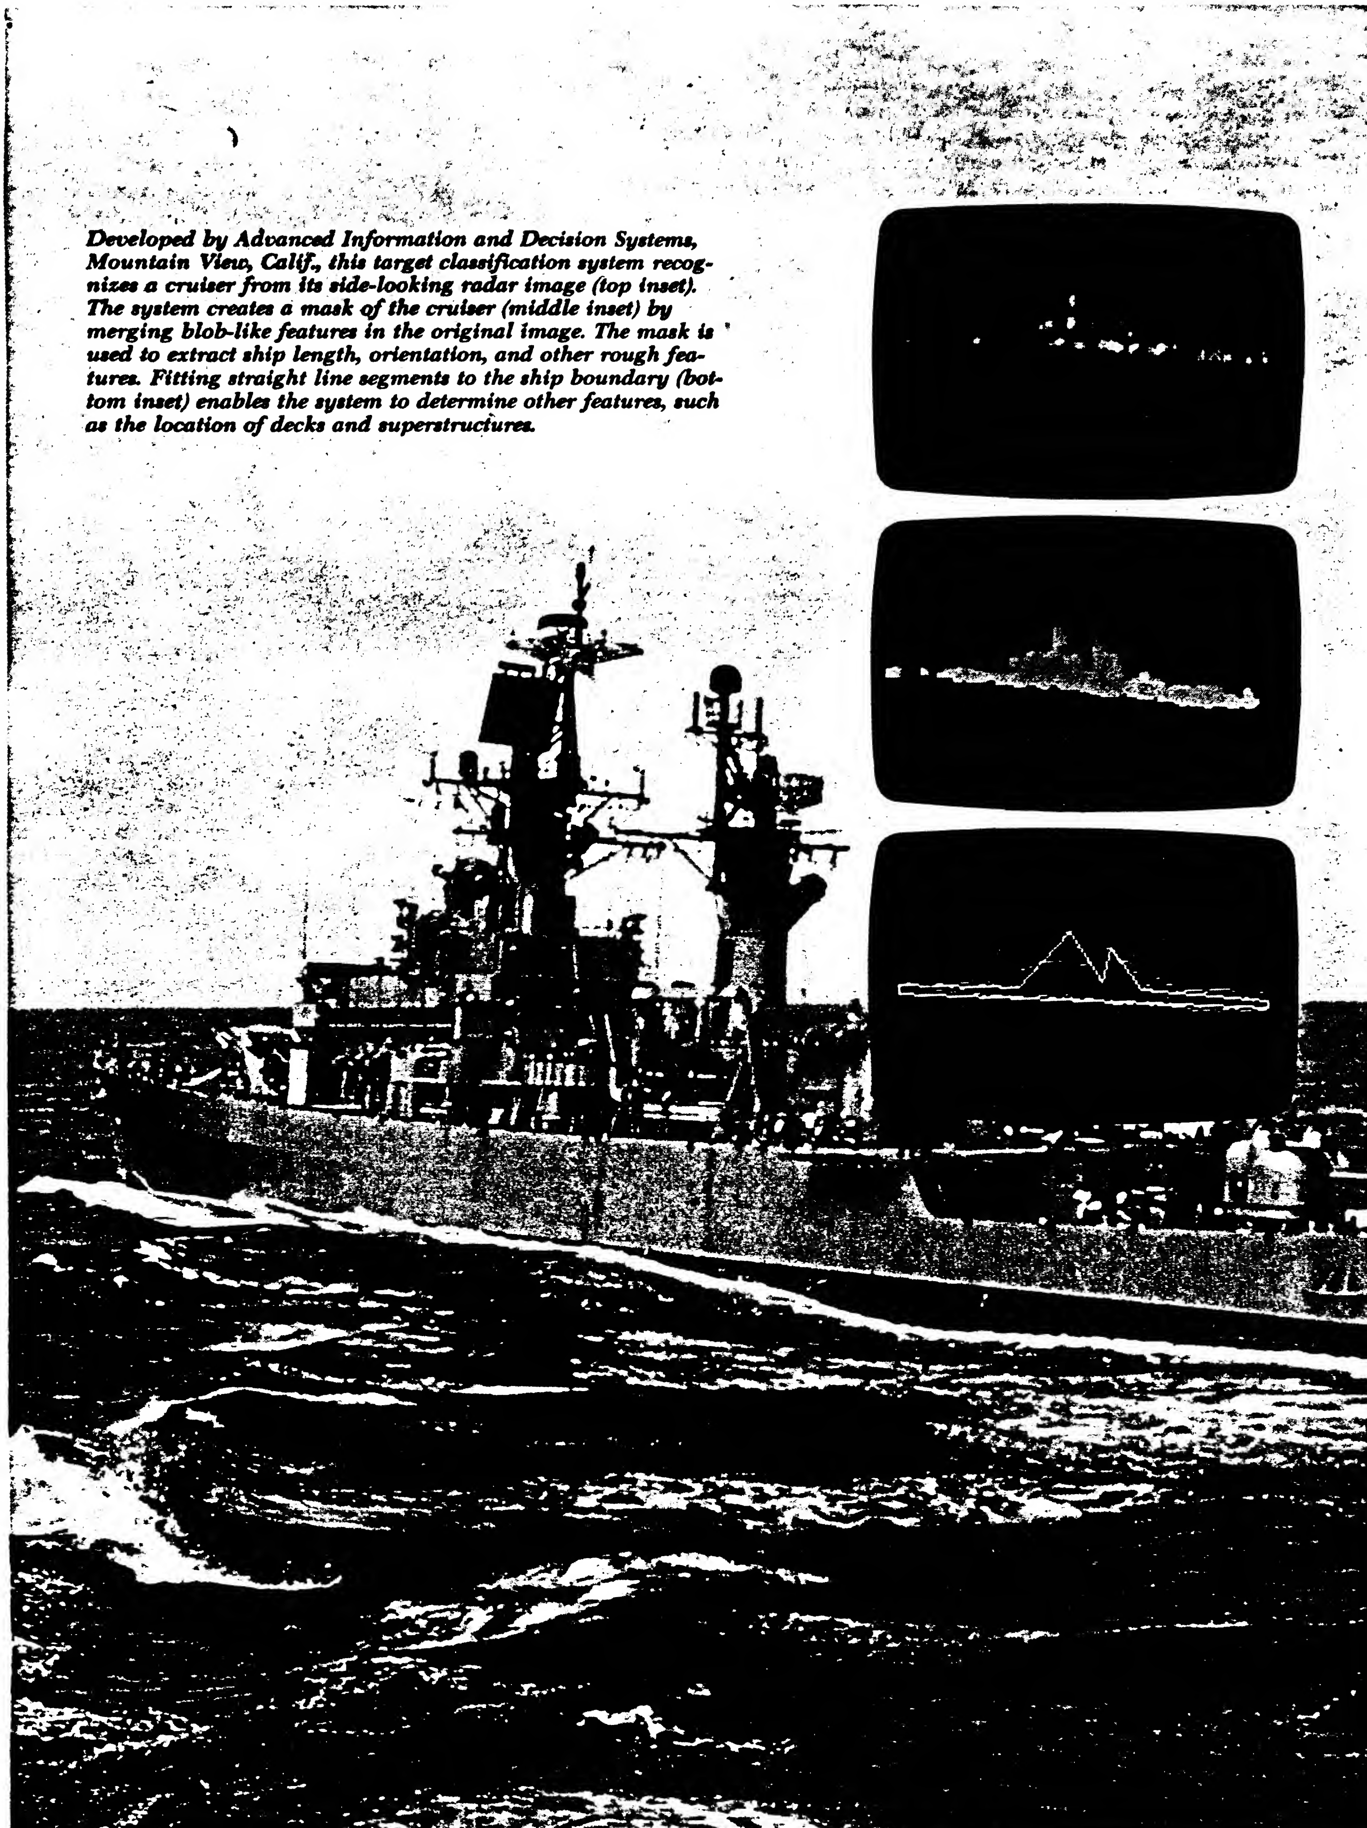

Developed by Advanced Information and Decision Systems, Mountain View, Calif., this target classification system recognizes a cruiser from its side-looking radar image (top inset). The system creates a mask of the cruiser (middle inset) by merging blob-like features in the original image. The mask is used to extract ship length, orientation, and other rough features. Fitting straight line segments to the ship boundary (bottom inset) enables the system to determine other features, such as the location of decks and superstructures.

US NAVY

fication is made.

A development that will hasten the spread of AI systems into commercial applications is the emergence of computers that efficiently run software written in LISP—the lingua franca of the AI community. LISP software runs on these machines, which sell for less than $100,000, as fast as on a mainframe computer selling for several million dollars. Xerox Corp. and two MIT spin-offs—Symbolics Inc. (Cambridge, Mass.) and LISP Machine Inc. (Culver City, Calif.)—currently sell LISP machines based on designs developed at Xerox PARC and MIT, respectively.

Some researchers are pressing ahead to develop hardware that will achieve high speed by using many individual processors working in concert like the neurons of the human brain. MIT, for example, is developing a parallel machine, built on an integrated circuit chip, that will solve a problem by breaking it into subproblems and distributing them among its processors. Another MIT chip will use parallel processors to speed the search through the branches of a semantic network—a common operation in AI systems. Such chips could lead the way to cheap—and incredibly powerful—thinking machines. □

For further information, see Resources *on p. 113.*

Paul Kinnucan is a senior editor of HIGH TECHNOLOGY.

BUSINESS OUTLOOK: Artificial intelligence

"There is no artificial intelligence industry . . . only artificial intelligence individuals," states Laura Conigliaro, security analyst at Bache. "You know that an industry has not yet developed when you find yourself concentrating on the technology instead of the companies," she says.

The AI field is too young for meaningful figures on market size, dollar volume, or growth trends. But major product groupings are developing that indicate the industry's future shape.

Expert systems and computer-aided instruction (CAI) appear to have the greatest market potential. These systems embody the accumulated knowledge of authorities—in fields like geologic prospecting or medical diagnosis.

Natural language programs that allow computer users to ask questions in ordinary English represent a core product for AI. Already programs for accessing large data bases are on the market.

Intelligent vision systems that permit robots to make judgments based on visual inputs are perhaps five years away. Industrial uses such as assembly, machining, and manipulating parts and equipment should offer a vast market when these systems attain more sophistication.

An additional market is for LISP machines, designed to execute directly programs written in LISP, a computer language ideally suited for developing artificial intelligence systems.

Expert systems. Programs that put advanced expertise in the hands of less-trained, lower-salaried workers are already appearing. Teknowledge (Palo Alto, Calif.) and Smart Systems, Inc. (Alexandria, Va.) are training programmers in the skills of knowledge engineers—those who interview authorities on specific subjects and put the knowledge into computer-assisted instruction and expert systems.

The expert system Explorer, developed by Cognitive Systems, Inc. of New Haven, Conn., replaces computer technicians who interface between geologists and a data base at oil companies. The market for this program is estimated at $10 million. Each of the larger oil companies has 5 to 10 major data bases which constitute additional markets.

Computer-aided instruction. Natural language systems will help revolutionize computer-assisted instruction. Programs under development enable students to ask questions of the computer and receive insightful advice, an educational experience comparable to one-on-one instruction with a human expert. Computer Thought Corp., Richardson, Tex., has developed a program (using LISP) to teach ADA, the new Department of Defense computer language, to computer programmers. A prototype system is expected by May 1983.

Programs for accessing large data bases are a prime application for natural language systems. After three decades of using computers, many large corporations had created massive data bases that they could no longer conveniently access. Artificial Intelligence Corp. (AIC), Waltham, Mass., developed its Intellect program to solve this problem for companies with IBM machines. With Intellect, the user can access the data base in ordinary English. The program enables the computer to understand and respond, thus establishing man–machine dialogue. Customers include Bank of America, Aetna Casualty, SoCal Edison, and Avco. Intellect has been on the market since 1981. AIC expects to have sold 200 systems at $70,000 each by January 1983.

Vision systems. AI research also focuses on improving optical character recognition to enable robots to make decisions based on visual input. Today, the computer can visually access data in only the most primitive way. The market potential is very large. Machine Intelligence Corp., Sunnyvale, Calif., Automatix, Inc., Burlington, Mass., Robotic Vision Systems, Melville, N.Y., and Xerox, Palo Alto, Calif., are engaged in this work.

LISP machines. Three companies manufacture special computers that efficiently run programs written in LISP. Xerox and the other two companies—LISP Machines, Inc., Culver City, Calif., and Symbolics Inc., Cambridge, Mass.—have sold a number of these machines at current prices between $50,000 and $144,000.

Right now, the AI world is composed of 15–20 small privately held companies that have been funded by venture capital firms located mainly in Boston and San Francisco. Although Wall Street is skeptical and uninvolved now, the venture capital firms believe in AI and continue to search enthusiastically for likely scientists to bankroll.

Some 6–12 large corporations, including IBM, Xerox, Schlumberger, and DEC, have AI research and development programs, but mainly to improve internal operations and products.

Many of the small AI companies were founded by university professors (mostly from Stanford and MIT) and the companies continue to have loose ties to academia. Consulting fees and seminars are an important source of income.

Because of their academic background, many AI people may feel more secure in research than product development and marketing. The ill-defined nature of the market helps rationalize delays in products; instead, firms concentrate on customized systems. Except for some machine vision systems, the Intellect program is the only off-the-shelf AI product on the market.

What will it take to get AI moving—with standardized commercial products? AIC executives say they succeeded by focusing on a specific business problem and being willing to commit the required funds—$10 million since 1975.

AI will someday become a full-fledged industry, but its growth will be evolutionary. We may not see the changes on a day-to-day basis but its overall impact in 10 or 20 years will be revolutionary. In the process, little-known small companies may emerge as giants.—***I.G.***